UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/804,945	03/19/2004	Curtis R. Priem	NVDA P000455	9182

26291 7590 01/13/2009

PATTERSON & SHERIDAN L.L.P. NJ Office
3040 Oak Post Road
Suite 1500
Houston, TX 77056-6582

EXAMINER	
CLEARY, THOMAS J	
ART UNIT	PAPER NUMBER
2111	

MAIL DATE	DELIVERY MODE
01/13/2009	PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

UNITED STATES PATENT AND TRADEMARK OFFICE

Commissioner for Patents
United States Patent and Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450
www.uspto.gov

BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

Application Number: 10/804,945
Filing Date: March 19, 2004
Appellant(s): PRIEM, CURTIS R.

Stephanie Winner
For Appellant

EXAMINER'S ANSWER

This is in response to the appeal brief filed 24 November 2008 appealing from the Office action mailed 6 March 2008.

(1) Real Party in Interest

A statement identifying by name the real party in interest is contained in the brief.

(2) Related Appeals and Interferences

The examiner is not aware of any related appeals, interferences, or judicial proceedings which will directly affect or be directly affected by or have a bearing on the Board's decision in the pending appeal.

(3) Status of Claims

The statement of the status of claims contained in the brief is correct.

(4) Status of Amendments After Final

The appellant's statement of the status of amendments after final rejection contained in the brief is incorrect.

The amendment after final rejection filed on 25 April 2008 has not been entered.

(5) Summary of Claimed Subject Matter

The summary of claimed subject matter contained in the brief is correct.

(6) Grounds of Rejection to be Reviewed on Appeal

The appellant's statement of the grounds of rejection to be reviewed on appeal is substantially correct. The changes are as follows: Claims 1-15, 17-18, and 23 stand

rejected as unpatentable over Ramakrishnan, Jones, and knowledge commonly known in the art, *as evidenced by* Browning.

(7) Claims Appendix

The copy of the appealed claims contained in the Appendix to the brief is correct.

5,826,081	ZOLNOWSKY	10-1998
6,633,897	BROWNING ET AL.	10-2003
5,812,844	JONES ET AL.	9-1998
6,085,215	RAMAKRISHNAN ET AL.	7-2000
2002/0083143	CHENG	6-2002

(9) Grounds of Rejection

The following ground(s) of rejection are applicable to the appealed claims:

Claims 1, 3-4, 9-10, and 23 are rejected under 35 U.S.C. 103(a) as being unpatentable over US Patent Number 5,826,081 to Zolnowsky ("Zolnowsky"), US Patent Number 6,633,897 to Browning et al. ("Browning") and US Patent Number 5,812,844 to Jones et al. ("Jones").

In reference to Claim 1, Zolnowsky discloses a method for scheduling the service of a thread, said method comprising the steps of: masking interrupts from hardware devices in order to ignore interrupts for other threads (See Column 6 Lines 34-42);

acquiring information associated with a thread (see Column 6 Lines 45-52); unmasking interrupts from the hardware devices in order to detect interrupts for the other threads (See Column 6 Lines 34-42); and rearranging an order in which the thread and the other threads will be serviced to schedule the thread for processing in accordance with said information, wherein the rearranging is performed simultaneously for the thread and the other threads (See Column 6 Lines 48-52). Zolnowsky further discloses that there are a plurality of dispatch queues, and for each single dispatch queue, a variety of different queuing mechanisms based on the information associated with the thread can be used for scheduling the threads associated with that queue depending on the real time application scheduling requirements (See Column 6 Lines 45-52). Zolnowsky does not disclose that the threads are ordered in a single queue corresponding to all requests received from the hardware devices. Zolnowsky further discloses that a real time operating system must be capable of scheduling a particular process within a fixed time limit (See Column 2 Lines 42-51). Zolnowsky does not explicitly disclose that the information regarding the thread is latency information and that the queuing mechanism (thread scheduler) used to sort each dispatch queue rearranges the order in which the threads will be serviced in a single queue based on said latency information. Browning discloses a system for scheduling thread execution from a single run queue instead of a global run queue and multiple local run queues (See Column 1 Line 22 – Column 2 Line 32). Jones discloses the use of deadline scheduling in which the order in which threads are scheduled for execution is simultaneously rearranged based on the latency of the thread (See Column 3 Lines 1-26).

It would have been obvious to one of ordinary skill in the art at the time the invention was made to construct the device of Zolnowsky using the single run queue of Browning and a latency based queuing mechanism, resulting in the invention of Claim 1, in order to optimize the thread processing (See Column 2 Lines 1-3 of Browning) and because Zolnowsky discloses that any queuing mechanism can be used (See Column 6 Lines 45-52 of Zolnowsky). Because both Zolnowsky and Jones disclose mechanisms for queuing threads (thread schedulers), it would have been obvious to substitute one queuing mechanism for the other to achieve the predictable result of scheduling the threads for servicing and execution.

In reference to Claim 3, Zolnowsky, Browning, and Jones disclose the limitations as applied to Claim 1 above. Jones further discloses computing the time at which the thread needs to be processed by summing the latency information with a current time (See Column 3 Lines 5-16).

In reference to Claim 4, Zolnowsky, Browning, and Jones disclose the limitations as applied to Claim 1 above. Jones further discloses that said latency information represents a time duration that is necessary to service the thread (See Column 3 Lines 5-16).

Claim 9 recites limitations which are substantially equivalent to those of Claim 1 and is rejected under similar reasoning.

Claim 10 recites limitations which are substantially equivalent to those of Claim 3 and is rejected under similar reasoning.

In reference to Claim 23, Zolnowsky, Browning, and Jones disclose the limitations as applied to Claim 1 above. Zolnowsky further discloses that the thread and at least one of the other threads correspond to interrupt requests from a single one of the hardware devices (See Column 1 Lines 31-39 and Column 7 Lines 44-49).

Claims 1-15, 17-18, and 23 are rejected under 35 U.S.C. 103(a) as being unpatentable over US Patent Number 6,085,215 to Ramakrishnan et al. ("Ramakrishnan"), Jones, and knowledge commonly known in the art, as evidenced by Browning and admitted to be prior art by the Appellant.

In reference to Claim 1, Ramakrishnan discloses a method for scheduling the service of a thread, said method comprising the steps of: masking interrupts from hardware devices in order to ignore interrupts for other threads (See Column 4 Lines 44-48, and Column 5 Lines 23-28); acquiring a latency information associated with a thread, wherein the latency information indicates a time at which the thread needs to be processed (See Column 10 Lines 48-64); and unmasking interrupts from the hardware devices in order to detect interrupts for the other threads (See Column 4 Lines 44-48, Column 5 Lines 23-28, and Column 7 Lines 22-52). Ramakrishnan further discloses the

thread and the other threads correspond to all requests received from the hardware devices (See Column 4 Lines 15-31), but is silent as to how the threads are stored. Official Notice is taken that it is notoriously old and well known in the art to store threads for processing is a single queue, as evidenced by Browning (See Column 1 Line 22 – Column 2 Line 32). This has been admitted by Appellant to be prior art. Ramakrishnan further discloses the use of deadline latency information when scheduling the threads (See Column 10 Lines 48-64), but does not disclose rearranging an order in which the thread and the other threads will be serviced in a single queue to schedule the thread for processing in accordance with said latency information, wherein the rearranging is performed simultaneously for the thread and the other threads. Jones discloses the use of deadline scheduling in which the order in which threads are scheduled for execution is simultaneously rearranged based on the latency of the thread (See Column 3 Lines 1-26).

It would have been obvious to one of ordinary skill in the art at the time the invention was made to construct the device of Ramakrishnan with the threads stored in a queue and with a latency based scheduling algorithm instead of a round robin scheduling algorithm, resulting in the invention of Claim 1, because Ramakrishnan is silent as to how the threads are stored and one of ordinary skill in the art would naturally look to various known methods of storing threads for processing, such as the use of a single queue which will optimize thread processing (See Column 2 Lines 1-3 of Browning), and because Ramakrishnan discloses that the latency requirements of the threads are important when scheduling threads (See Column 10 Lines 48-64 of

Ramakrishnan. Because both Ramakrishnan and Jones disclose algorithms for queuing threads (thread schedulers), it would have been obvious to substitute one queuing algorithm for the other to achieve the predictable result of scheduling the threads for servicing and execution.

In reference to Claim 2, Ramakrishnan, knowledge commonly known in the art as evidenced by Browning, and Jones disclose the limitations as applied to Claim 1 above. Ramakrishnan further discloses that said latency information is computed based on a buffer size (See Column 11 Lines 24-47).

In reference to Claim 3, Ramakrishnan, knowledge commonly known in the art as evidenced by Browning, and Jones disclose the limitations as applied to Claim 1 above. Jones further discloses computing the time at which the thread needs to be processed by summing the latency information with a current time (See Column 3 Lines 5-16).

In reference to Claim 4, Ramakrishnan, knowledge commonly known in the art as evidenced by Browning, and Jones disclose the limitations as applied to Claim 1 above. Jones further discloses that said latency information represents a time duration that is necessary to service the thread (See Column 3 Lines 5-16).

In reference to Claim 5, Ramakrishnan, knowledge commonly known in the art as evidenced by Browning, and Jones disclose the limitations as applied to Claim 1 above.

Ramakrishnan further discloses that said latency information represents a maximum time allowed before a first buffer will be emptied and a read operation will switch to process a second buffer (See Column 11 Lines 24-47).

In reference to Claim 6, Ramakrishnan, knowledge commonly known in the art as evidenced by Browning, and Jones disclose the limitations as applied to Claim 1 above. Ramakrishnan further discloses that said latency information represents a time duration that is necessary to setup the thread to perform interrupt processing for the thread (See Column 11 Lines 24-47).

In reference to Claim 7, Ramakrishnan, knowledge commonly known in the art as evidenced by Browning, and Jones disclose the limitations as applied to Claim 1 above. Ramakrishnan further discloses that said latency information is dependant on a hardware constraint for one of the hardware devices (See Column 11 Lines 24-47 and Column 12 Lines 43-55).

In reference to Claim 8, Ramakrishnan, knowledge commonly known in the art as evidenced by Browning, and Jones disclose the limitations as applied to Claim 1 above. Ramakrishnan further discloses that said latency information is provided by a device driver (See Column 12 Lines 43-55).

Claim 9 recites limitations which are substantially equivalent to those of Claim 2 and is rejected under similar reasoning.

Claim 10 recites limitations which are substantially equivalent to those of Claim 3 and is rejected under similar reasoning.

Claim 11 recites limitations which are substantially equivalent to those of Claim 7 and is rejected under similar reasoning.

In reference to Claim 12, Ramakrishnan, knowledge commonly known in the art as evidenced by Browning, and Jones disclose the limitations as applied to Claim 11 above. Ramakrishnan further discloses that said hardware constraint is a size of a buffer (See Column 11 Lines 24-47).

In reference to Claim 13, Ramakrishnan, knowledge commonly known in the art as evidenced by Browning, and Jones disclose the limitations as applied to Claim 11 above. Ramakrishnan further discloses that said hardware constraint is a fullness of a buffer (See Column 11 Lines 24-47).

In reference to Claim 14, Ramakrishnan, knowledge commonly known in the art as evidenced by Browning, and Jones disclose the limitations as applied to Claim 11

above. Ramakrishnan further discloses that said hardware constraint is dynamically computed based on a buffer size (See Column 11 Lines 41-42).

Claim 15 recites limitations which are substantially equivalent to those of Claim 8 and is rejected under similar reasoning.

In reference to Claim 17, Ramakrishnan, knowledge commonly known in the art as evidenced by Browning, and Jones disclose the limitations as applied to Claim 1 above. Ramakrishnan further discloses toggling an interrupt line (See Column 10 Lines 31-47).

In reference to Claim 18, Ramakrishnan, knowledge commonly known in the art as evidenced by Browning, and Jones disclose the limitations as applied to Claim 1 above. Ramakrishnan further discloses determining the thread should be activated; and activating the thread for processing (See Column 4 Lines 16-32).

In reference to Claim 23, Ramakrishnan, knowledge commonly known in the art as evidenced by Browning, and Jones disclose the limitations as applied to Claim 1 above. Ramakrishnan further discloses that the thread and at least one of the other threads correspond to interrupt requests from a single one of the hardware devices (See Column 4 Line 15 – Column 5 Line 3 and Column 9 Lines 23-27).

Claims 21 and 22 are rejected under 35 U.S.C. 103(a) as being unpatentable over Ramakrishnan, knowledge commonly known in the art as evidenced by Browning, and Jones as applied to Claim 1 above, and further in view of US Patent Application Publication Number 2002/0083143 to Cheng ("Cheng").

In reference to Claims 21 and 22, Ramakrishnan, knowledge commonly known in the art as evidenced by Browning, and Jones disclose the limitations as applied to Claim 1 above. Ramakrishnan further discloses creating threads for interrupt processing (See Column 10 Lines 34-37). Ramakrishnan, knowledge commonly known in the art as evidenced by Browning, and Jones do not disclose creating the thread for interrupt processing when one of the hardware devices is initialized, wherein the thread is created for use during processing of a first interrupt that the one or more hardware devices is configured to generate; and freeing the thread when the one of the one or more hardware devices is shut down, as in Claim 21, and creating an additional thread, wherein a first interrupt identification number is associated with the thread and a second interrupt identification number that is different than the first interrupt identification number is associated with the additional thread and the additional thread is created for use during processing of a second interrupt that the one of the hardware devices is configured to generate; and freeing the additional thread when the one of the one or more hardware devices is shut down, as in Claim 22. Cheng discloses that it is well known to create a thread when a device is added to a system and to free a thread when a device is removed from a system (See Figure 6 and Paragraphs 67-73 and 24). The

device of Cheng would inherently use different identification numbers for the thread, as the device would be inoperable the same identification number was used for multiple interrupt threads.

It would have been obvious to one of ordinary skill in the art at the time the invention was made to use the thread creation of Cheng in the device of Ramakrishnan, Browning, and Jones, resulting in the invention of Claims 21 and 22, because Ramakrishnan is silent as to how the threads are created and one of ordinary skill in the art would naturally look to methods of creating threads; and to allow both plug and play and non plug and play devices to be used in the same network system (See Abstract and Paragraphs 8 and 74).

(10) Response to Argument

Official Notice was taken in the previous Office Action. To adequately traverse such a finding, an Appellant must specifically point out the supposed errors in the Examiner's action, which would include stating why the noticed fact is not considered to be common knowledge or well-known in the art. If the Appellant does not traverse the Examiner's assertion of Official Notice, the common knowledge or well-known in the art statement is taken to be admitted prior art. (See MPEP 2144.03 C). The Examiner's assertion of Official Notice is hereby taken to be admitted prior art due to the Applicant's failure to traverse the assertion.

Appellant has argued that Zolnowsky does not disclose that all of the requests are ordered in a single queue, as Zolnowsky discloses the use of multiple queues (See Page 10 Paragraph 3). In response, the Examiner notes that, as indicated in the rejections, Zolnowsky was not relied upon to disclose this limitation.

Appellant has argued that Zolnowsky does not disclose masking and unmasking interrupts from hardware devices (See Page 10 Paragraph 4 – Page 11 Paragraph 1). In response, the Examiner notes that the broadest reasonable interpretation of masking and unmasking interrupts is blocking interrupts from being sent to a processor (masking) and allowing interrupts to be sent to a processor (unmasking). The queue lock of Zolnowsky prevents any devices not having the lock from performing operations on the queue (See Column 6 Lines 34-38). Zolnowsky further discloses that interrupts of a processor are achieved through the use of interrupt threads (See Column 8 Lines 55-64). Thus, when the queue lock is acquired, interrupts are prevented from being sent to the other processors, and are thus masked. When the queue lock is released, interrupts can one again be sent to the other processors, and are thus unmasked. The Examiner further notes that the portion of Zolnowsky referred to by Appellant on Page 11 Paragraph 1 is found out Column 8 Lines 63-64, and not Column 8 Lines 3-4.

Appellant has argued that Jones does not disclose ordering all requests from all of the hardware devices in a single queue, nor that interrupts are masked or unmasked

(See Page 11 Paragraph 2). In response, the Examiner notes that, as indicated in the rejections, Jones was not relied upon to disclose these limitations.

Appellant has argued that Browning does not disclose ordering all requests from all of the hardware devices in a single queue, as Browning describes the preferred embodiment of the run queue as being "subdivided into 128 first-in-first-out (FIFO) queues, where there is a unique queue for each priority level" (See Page 11 Paragraph 3 – Page 12 Paragraph 1). In response, the Examiner notes that Browning discloses that this is a preferred embodiment (See Column 5 Lines 6-11), but is not limited as such (See Column 5 Line 66 - Column 6 Line 3). Browning clearly discloses the use of a single run/execution queue (See Column 1 Lines 54-56, Column 3 Lines 28-35, and Abstract) and does not require that the single queue have sub-queues, nor prohibit the single queue from having no sub-queues (See Column 3 Lines 28-51). Furthermore, even when using the preferred embodiment, the sub-queues are still, as a whole, a single run queue ("*The run queue* is subdivided…").

Appellant has argued that neither Zolnowsky, Browning, nor Jones disclose that all of the threads are simultaneously rearranged in a single queue (See Page 12 Paragraph 2). In response, the Examiner notes that Browning discloses a system for scheduling thread execution from a single run queue instead of a global run queue and multiple local run queues (See Column 1 Line 22 – Column 2 Line 32); and Jones discloses the use of deadline scheduling in which the order in which threads are

scheduled for execution is simultaneously rearranged based on the latency of the thread (See Column 3 Lines 1-26).

Appellant has argued that Ramakrishnan does not disclose ordering all requests from all of the hardware devices in a single queue and that all of the threads are simultaneously rearranged in the single queue (See Page 12 Paragraph 4). In response, the Examiner notes that, as indicated in the rejections, Ramakrishnan was not relied upon to disclose these limitations.

Appellant has argued that Ramakrishnan does not disclose that the threads are ordered according to priority (See Page 12 Paragraph 4 – Page 13 Paragraph 1). In response, the Examiner notes that, as indicated in the rejections, Ramakrishnan was not relied upon to disclose this limitation. Ramakrishnan discloses the thread and the other threads correspond to all requests received from the hardware devices (See Column 4 Lines 15-31), but is silent as to how the threads are stored. Official Notice was taken that it is notoriously old and well known in the art to store threads for processing is a single queue, as evidenced by Browning (See Column 1 Line 22 – Column 2 Line 32). This has been admitted by Appellant to be prior art. Ramakrishnan further discloses the use of deadline latency information when scheduling the threads (See Column 10 Lines 48-64), but does not disclose rearranging an order in which the thread and the other threads will be serviced in a single queue to schedule the thread for processing in accordance with said latency information, wherein the rearranging is

performed simultaneously for the thread and the other threads. Jones discloses the use of deadline scheduling in which the order in which threads are scheduled for execution is simultaneously rearranged based on the latency of the thread (See Column 3 Lines 1-26). Thus, Ramakrishan in combination with knowledge commonly known in the art and Jones discloses the limitation.

Appellant's arguments with respect to Jones and Browning in combination with Ramakrishnan are substantially equivalent to those presented with respect to their combination with Zolnowsky (See Page 13 Paragraph 2). The Examiner's response is therefore the same. The Examiner notes that the rejections involving Ramakrishnan do not rely on a combination with Browning, but rather on a combination with what was well known in the art, which Browning was relied upon to provide evidence of. As indicated above, Appellant's failure to traverse the Examiner's assertion of Official Notice has been taken as an admission of prior art. Further, as indicated in the rejections, in the combinations involving Zolnowsky, Browning was relied upon to disclose a system for scheduling thread execution from a single run queue instead of a global run queue and multiple local run queues; while in the combinations involving Ramakrishnan, Browning was relied upon to provide evidence of Official Notice taken by the Examiner that it is notoriously old and well known in the art to store threads for processing is a single queue.

Appellant has argued that Cheng does not disclose processing interrupts (See Page 13 Paragraph 5 – Page 14 Paragraph 6). In response, the Examiner notes that Ramakrishnan was relied upon to disclose creating threads for interrupt processing (See Column 10 Lines 34-37). Cheng was relied upon to disclose that it is well known to create a thread when a device is added to a system and to free a thread when a device is removed from a system (See Figure 6 and Paragraphs 67-73 and 24). The device of Cheng would inherently use different identification numbers for the thread, as the device would be inoperable the same identification number was used for multiple interrupt threads. One cannot show nonobviousness by attacking references individually where the rejections are based on combinations of references. See *In re Keller*, 642 F.2d 413, 208 USPQ 871 (CCPA 1981); *In re Merck & Co.,* 800 F.2d 1091, 231 USPQ 375 (Fed. Cir. 1986).

(11) Related Proceeding(s) Appendix

No decision rendered by a court or the Board is identified by the examiner in the Related Appeals and Interferences section of this examiner's answer.

For the above reasons, it is believed that the rejections should be sustained.

Respectfully submitted,

/Thomas J. Cleary/

Patent Examiner, Art Unit 2111

Conferees:

/Paul R. Myers/
Primary Examiner, Art Unit 2111

Mark Rinehart
/M. R./
Supervisory Patent Examiner, Art Unit 2111